

## AWIPS II Technical Interchange Meeting

Primary audience: GOES-R Proving Ground

Time: 2 PM to 3:30 PM EST

Date: January 12, 2011

An audio recording of the call is available.

Notes prepared by Jordan Gerth

Attendance: Jordan Gerth, Kaba Bah, Hiro Gosden, Ed Mandel, Deb Molenaar, Brian Motta, Steve Miller, Eric Lau, Gregg Grosshans, NASA SPoRT (including Matt Smith, Kevin McGrath, and Jason Burks), Bill Campbell, Doug Lawson and Frank Griffith from Raytheon

Introduction: The main thrust of today is to go through the action items. Other topics can be introduced at the end of the call.

- **Action: Work to collect, and then distribute information on the McIDAS AREA plugins currently developed by Raytheon, NCEP, and SPoRT so that a unified decoder can be delivered operationally to support the experimental product ingest.**
  - **OPR: Raytheon, NCEP and SPoRT to provide code or pseudo-code so that the group can decide if this is worth pursuing. (Nov 12)**

NCEP has delivered their version of the McIDAS AREA decoder (for ingest on the edex side) to Raytheon Technical Services (RTS) and the display end was to be delivered today.

Raytheon noted that their plug-in was a prototype done a few years ago and not at the level of the NCEP development which will be made available to us shortly. The timeframe is unclear. David Plummer previously said that ingest, storage, and edex server side handling of the McIDAS AREA data was delivered in October. Display capabilities (viz) were delivered to Raytheon very recently. Some of the NCEP code has been integrated. Raytheon says this is a matter of entering it into the AWIPS II discrepancy report (DR) process. Bringing it into the DR process means it is being rationalized among all of the other DRs being done right now. Raytheon will probably support it once it is in the baseline (for maintenance changes, etc.). There is some lack of clarity in this area. If that was too far in the future, it may be expedited through some kind of DR. We are thinking about the March or April time frame.

*Update to action: Ed will determine when it will be in the baseline.*

We previously discussed the capability for the 24-bit display. One of things that happened in the interim was a new, unlimited overlay feature in AWIPS II which was not in AWIPS I. It was approved in the summer of 2010. When it comes down to 24-bit display, it has to do with how the McIDAS format is displayed. The other satellite format is still confined to 8 bits. Having a

good timeframe for when we will see the code will help in investigating what additional work is needed to maximize the capability of a single format.

We can do a 24-bit display. Jordan said he was investigating the end-to-end handling of some of our POES imagery into the software and looking closely at how the `SatBlendedResource` and `SatBlendedResourceData` works. It really comes down to the alpha channel in the brightness of the image. Also, we need to consider what may need to be done to make that more accessible. It would be nice to have a 0 to 255 red, blue, and green color bar, and then assign it to each satellite product with the click of a button. It is the RGB that we desire, even though we are taking three 8-bit products. The user has to go point by point, load a product, and then change the color enhancement. Essentially, you can load up three things and it turns into a RGB.

Frank says there are a couple ways of doing it. One way is with the Java code itself. A more interesting way of doing it would be to load the three pieces into the graphics card and let the graphics card do the manipulation with color bars.

Steve Miller interjected that some of the products would not benefit from the user being able to manipulate them in this fashion. He suggests loading and combining different images as a prior process to the display. In some products, that is a good thing to do. In others it may not be.

Some of the xml files where they call the `SatBlendedResource` it may be possible to put three different bands in there and it would do what is desired.

There are also a series of files which perform image manipulations. For example, there is a Python script for water vapor divided by IR window data. A Python code does the blending and creates the product for display. That is closer to what you want to do. It takes multiple inputs.

Image subtractions are working with 8-bit space. The other thing was exactly how time matching is done. Files are required to have the same time in order to blend them. It is possible to relax the time matching rules. For scientific computations, you may want to relax them somewhat. For GOES, they wanted for exactly the same satellite at exactly the same time.

It does not have to be 8-bit. The mechanism should scale in terms of bytes. As far as we know, the record format has the number of bytes for the depth of the record. Most of it is one byte. Gridded data is 32-bit float, which is loaded into the graphics card.

Image compositing is handled in:

- `com.raytheon.viz.satellite/src/com/Raytheon/viz/satellite/rsc/SatBlendedResource.java`
- `com.raytheon.viz.satellite/src/com/Raytheon/viz/satellite/rsc/SatBlendedResourceData.java`

Image subtraction is handled in:

- `build/static/common/cave/etc/derivParamScripts/poesDif11u3_7uIR.py`

- build/static/common/cave/etc/derivParamScripts/satDif11u12uIR.py
- build/static/common/cave/etc/derivParamScripts/satDif11u3\_9uIR.py
- build/static/common/cave/etc/derivParamScripts/satDivWVIR.py

Deb Molenaar inquired about the three separate files approach and how they are somehow combined on the fly or if you would be able to send it in one file scaled. Using the current software, you can send three separate 8-bit arrays. The data comes one record at a time, each record as a separate file.

You can have a Python script do the combining and making a composite.

Jason at SPoRT commented that they used the AREA decoder. It was able to display a different range of projections. They populated the database. We can decode the full suite of McIDAS products. NCEP was working on their decoder. They have been working on some other plug-ins in the interim and have left that for now (not actively developing it). They are certainly interested in seeing what NCEP has. They have been successful in visualization.

*Update to action: Ed Mandel will chase down when the McIDAS decoder will be available to the GRPG partners.*

Deb asked what else SPoRT had done. They had data from the lightning mapping array in netCDF format and smoke data in kml format from the hazardous mapping service. Each of those has an edex plug-in for decoding as well as a viz plug-in for visualization. They can display a single point for fire hot spots. They also have gridded data they can display. Lastly, they have used their McIDAS plug-in.

- **Action: Raytheon needs to further describe how the legends are created and displayed. Information on the shader language and process is also important. The maximum and minimum values are parameterized when the data selection occurs, and that controls what you see on the screen as far as color.**
  - **OPR: Raytheon to provide comments on how their code works for the display legends. (Nov 12)**

In general, on the display side, there is SatelliteResource Java code. It builds the legend. It does this by grabbing data out of the satellite record metadata. It is mainly the time and the physical element that it uses. Then there is another file that maps the physical element to a human readable name (SatelliteConstant). If it does not have a physical element, and it does not map, then it will display the physical element as it is in the metadatabase. You can see all that by looking at the SatelliteResource code.

In SatelliteConstant, there are bunch of private static final strings which outline how the product identification in that file. Is that to say that, in order to get the product legend, we would need to add it or override? If you look at the SatelliteResource code itself, it will just display the metadata itself. It may be slightly hieroglyphic. The question was asked about why the legends

are hard-coded into the edex server code. Why is that not where the style information is? Raytheon says it might have been historical. This was one of the first plug-ins that was written. The style stuff evolved a little later. It would make sense to put all of that in the style xml file. In order to submit that for a design discussion, we must enter a TTR.

*Update to action: Draft a TTR. The TTR would make it flexible to make changes as to how the labels are displayed. When you are dealing with new satellite data, it will be easier and it is probably a worthwhile thing to change. At this point, we have to go through the TTR/DR process, however. This action item remains open.*

- **Tell us at what point in the ingest and display process is the conversion between bit space and value space performed. Does/could the interpolation done on the graphics card lead to a false return value?**
  - **OPR: Raytheon to provide this information at the next TIM at Raytheon's Omaha facility.**

There is nothing done on the ingest side specifically with the data. Data is stored however it comes in (retains original values). In most cases, it goes to the graphics card that way. That is where the style rules come into play and give you a mapping mechanism (engineering type value).

Examine the awips.satellite table within the metadatabase. The units are stored in the metadata. For satellite imagery, this comes from the GINI format. The decoder may somehow ascertain the units. It is not certain which pieces of metadata are used. We cannot determine what from the metadata are actually used for storage and display and there is really no document. The easiest way to answer that is to look at the fields which become part of the data URI, a subset of the total, which is used for discovery.

For display, the other interesting thing is that the majority (all except for the low cloud product) do not have a corner point. It goes to the satellite spatial table and uses the gid. In the postgres database, it stores a geometry record (GIS part of it). The coverage is created and stored in there. That is something across the board for all of the plug-ins. It is a consistent way of doing it across all of the data types. That includes the projection, space between the pixels, dx, dy, nx, ny, etc. The geometry is a very complete object.

*Update to action: This action remains open for an in-person TIM.*

- **How do we ensure the transaction of ingesting products into the data repository is complete every time? More broadly, we need to identify fail points in the software. Memory management and execution time should be closely watched to not impact the rest of the system.**
  - **OPR: Raytheon at next TIM.**

There is not a formal mechanism in the system for assuring a transaction is complete. It is possible to turn on SAtransaction, but that would be a big hit performance-wise. There are some lower-level ways of accommodating a complete transaction check. If a file fails somewhere, you leave it where it is so it is possible to try again. There is nothing equivalent to the AWIPS I file handling in AWIPS II in terms of ability to track bad data.

Is there anything that would get us to the point before the software is released? The logging indicates if something fails. The data store would still have the data file. The data comes off the SBN and into the LDM. That deposits it into the data store. It would be possible to fight bad data through going from the error messages to the data store. It is a little more awkward workaround. At one point, they were going to archive the raw data. There is an intention to keep a certain amount available. That would be the idea. I do not know how long.

Gregg noted that WDTB had mentioned they were going to archive the HDF5 and then dump the tables. WDTB was going to propose to break those files down on the URI side by time and radar site instead of putting them down into one bin. Are those TTRs in one system? There were a series of DRs generated for the enhanced capability.

*Update to action: Ed Mandel said he would check the DRs to see if there is a proposal to change the satellite storage. He believes this is on the list for future satellite releases TBD by SREC. This action remains open to report problems completing transactions.*

- **More investigation is needed to determine if the JAI libraries could be problematic with really large datasets.**
  - **OPR: Raytheon at next TIM**

With regards to the JAI, we talked about limitations and file size limitations with really large data sets. There has not been too much done to determine if there is an upper quantitative limit.

In R1G1-7, with respect to the HDF5 bins, there are now multiple images for the same scan and band. With the visible channel and the GOES data that comes in, that is at 1-km resolution. The data gets decimated five times. It is a quad decimation mechanism which is done by JAI. The HDF5 stores each of those decimations. You are grabbing different resolutions of the data. HDF5 chunking is used which creates the concept of tiles. Each array in the data store is retrieved by tiles. It covers the display area plus a little more. You are grabbing the right tiles and the right decimation out of the HDF5 file, which creates the Google experience for that kind of data, but there are limitations.

The Sun implementation of JAI is used for Linux and Windows. They decode behind the scenes and perform better than those which use JVM (a virtual machine). Have a layer of tiling before it gets to the JAI code has not been accommodated yet. Tests indicate that the software can deal with the MODIS data as-is. Once we deal with data that is larger than that, we will have to start worrying about it.

We decided that the maximum spatial resolution we desire is 250 meters. SPoRT has displayed 250 m. It had a tough time trying to load more than just a few images at the National scale. It was a hybrid product which was created by oversampling GOES to 250 m and inserting a MODIS swath into it, but that was a while ago when looping was difficult.

*Update to action: SPoRT has agreed to attempt to ingest and display 250 m resolution data into the latest release of AWIPS "II" prior to the next TIM. Action remains open.*

- Work with Raytheon to determine how the shader language works with the graphics card for displaying satellite imagery at greater than 8 bits, as well as compositing multiple images.
  - OPR: **Raytheon at next TIM.**

Frank said that basically the shader language has been pulled out into a separate file so it is easy to find. It is a file with an extension of glsl. It is called color map. All it really does is map the color table, which has been stored in graphics card as a one-dimensional texture. It maps back to the image pixel value and then there is a little bit of logic in there. It also takes the input from the slider bars for the brightness and contrast and controls the creation of the color bar that is in the corner of the viz window.

Prior to invoking the shader language, the style rules have been applied. The image has been dealt with by the style rules. The style rules are specific for each data element. It specified whether the color table is linear or logarithmic. It is useful to look at the SatelliteResource to see how the colors get applied.

There is no other documentation regarding shader language. You have to get down to reading the code and follow it through. There is not that much involved so it is not too bad.

SatelliteResource.java is the overall manager of how the display works. It builds the legend and grabs the style rules. It has the sampling code (when you sample the data).

*Update to action: GRPG partners should evaluate code based on this direction to determine if there are further questions. If not, this action can be closed.*

- **Beyond the current capabilities, discussing with Raytheon the exact process in linking colors to value mappings and defining a maximum and minimum for the displayed range. Can more of this be defined in the configurable xml?**
  - OPR: **Raytheon at next TIM.**

Visible satellite imagery is straightforward on an 8-bit scale from 0 to 255. Some of the other channels are more complicated. The one thing that is a little vexing is we now go more toward the formulaic approach instead of a piecewise definition. In AWIPS I, we used to set the color bars by defining the stored data value to a display value. Now, there are style rules applied to the resource. The style rules do not seem to have anything with regards to how the mapping occurs,

though they do contain a maximum and minimum value bar and control whether or not the displayed color range should be linear and logarithmic. If you look at the IR satellite, it is setup with a linear section then a bunch of discrete colors. That ends up building the xml color table.

Sometimes the data will be converted from bit space to value space and then color mapping happens in value space once it is in the graphics card. There is a lot of manipulation going on before it is loaded.

*Update to action: This action remains open.*

- **Determine what from the metadatabase is actually used for discovery and display.**
  - **OPR: Raytheon at next TIM.**

The HDF5 appears to store the data URI in group-type architecture. The metadatabase contains everything else and we really do not know what in the database is used.

The most important items are the ones that are in the data URI. It represents the hierarchical structure of the data URI. If you open up a HDF5 file, you will see the data URI. That is the field that would be used for discovery purpose.

Related file: <ftp://ftp.ssec.wisc.edu/pub/jordang/pg/h5dump-satellite-2010120322.txt>

You can bring up the product browser in the CAVE. The data used for discovery will appear in there. There are other metadata items that are used for the geolocation operations. Examine the spatial column and then trace the items through. A number of them are not used. If you search the code, there are no references to all of them.

Related file: <ftp://ftp.ssec.wisc.edu/pub/jordang/pg/psql-metadata-awips-satellite.txt>

Related file: <ftp://ftp.ssec.wisc.edu/pub/jordang/pg/psql-metadata-awips-satellite-spatial.txt>

*Update to action: This action remains open.*

- **Based on our future expectations from NPP and GOES-R, identify what needs to be handled in the metadatabase and what needs to be stored in the HDF5.**
  - **OPR: GOES-R PG Partners need to answer. (Nov 30)**

*Update to action: This action remains open. Jordan noted that the current performance is probably fine the way it is. GRPG partners should examine the related files above to understand the architecture of the metadatabase and the data bin then provide comments if necessary.*

- **Sample project: Using the display side to combine polar and geostationary data dynamically to create a maximum coverage, maximum resolution product as part of the display. Raytheon says this is probably within the range of what we can do by relaxing the time matching, or using a threshold in the time matching.**

- **OPR: SPoRT, NCEP, CIRA, CIMSS develop a sample AWIPS II development project that can be executed during the TIM at Raytheon's Omaha facility.**

Trying to blend different products remains a desirable project to pursue. We are searching for data sets which you will write into HDF5 that may be unique. If you have a unique dataset type which we can broaden to apply to other things within the Proving Ground, that would be most sought. There was some discussion about combining different satellite images or products into a single display.

If all of the satellite data will reside in a single HDF5 file, Steve Miller suggests that we look into investigating how to search through a HDF5 file. Once we get the AREA code delivered to us, we should have a better foot forward.

What is the latest field deployment date for this software? This is the date we need to target to deliver our compatible products, many of which are in netCDF. Ed said that the current schedule has the completion of field OT&E on September 30, 2011. At least, we are thinking the first quarterly release will be in the winter of 2012 (January). The GRPG partners have all of 2011 to get ready. Part of the Software Recommendation and Evaluation Committee (SREC) process will be staging enhancements incrementally.

For governance, we have some stuff already. Raytheon has done some work on this as well. We have a draft software development plan which has been used within SEC for the new development folks. There are lots of other pieces of governance, such as deployment.

*Update to action: This action remains open. Projects may still be submitted for a future TIM to be held in person.*

- **Developers should talk about what is done now with image blending. Subsequent TIMs should address capabilities of the graphics card.**
  - **OPR: Raytheon at next TIM.**

Brian noted that a possible small project could be to look at a large domain image composite and look at the spatial and temporal thresholds. Frank talked about relaxing those. Also, more information on how and when mosaicking is done, how that routine works, how it decides what pixels from which satellite to put where would be good to know. We could start with the current routine, with the combination of GOES-East and GOES-West imagery.

*Update to action: Consolidate with other actions open which are similar.*

- **Outline a strategy of how compositing and image subtractions should desirably work.**
  - **OPR: GOES-R Proving Ground Partners (Nov 30)**



Jordan has worked with image subtractions. Steve Miller asked about the ability to do loops of satellite imagery. GeoColor is a composite product pre-processed as such. GeoColor compositing will be done prior to deliver. We are talking about multiple data sources of the same type, such as blended TPW and microwave TPW. SPoRT is putting a swath of MODIS on top of another GOES image. That is a layering technique which can be done outside of AWIPS but may be able to be done within AWIPS, such as taking data values from your image and assign transparency vectors to them, creating a dynamic transparency and alpha channel. For a future capability suggestion, a dynamic transparency would be something worthwhile to submit.

There is a little bit of difference between image subtraction and image transparency. Transparency has to be done with blended resources and is handled in a different fashion compared to image subtraction, which is done with a Python script.

*Update to action: Raytheon and the GRPG partners should continue to discuss this at a subsequent TIM based on software developments and GRPG requirements.*

- **How does the satellite decoder work? What are the specific class-level interactions within the plug-in. Developer involvement is required on how to implement a new or generalized decoder, and identify any interference with other classes which may result.**
  - **OPR: Raytheon at next TIM.**

You probably do not want to modify the satellite decoder that is in there now. That deals with GINI. McIDAS plug-in would be independent of that. I would assume the McIDAS plug-in has its own HDF5 schema and metadata schema. It is completely independent of the GINI satellite decoder. Each plug-in has a mechanism. It is up to the display resource to know where to go and get the data.

*Update to action: Move this action item to the in-person TIM. At some point, we have to make a decision as to whether we want to create an experimental satellite decoder.*

Jordan asked Raytheon: “Is there any problem with putting the experimental data products in the same bin as the other products?” Not really, because there is a separate metadata record pointing to it. It would have a unique data URI which operationally no one cares about. It is often done now when people want to grab data. They run them in on the manual end point. It will go in based on the plug-in that is already there for that data type. It will go into the HDF5 file and does not usually cause any problems.

- **Acquiring, building, and developing in 64-bit at Proving Ground provider sites.**
  - **OPR: Ed Mandel said that whoever gets AWIPS II releases will find this capability as part of the February 2011 release.**

Ed clarified that they are putting the infrastructure to be able to do that but not necessarily turning it all on. There have been some discussion with Raytheon and a white paper put out.

*Update to action: Ed Mandel will have more information to share after he follows up.*

- **We need to find out the specifics of ingesting a local model into AWIPS II.**
  - **OPR: Raytheon at next TIM.**

Is there document available for ingesting a local model in GRIB2 format into AWIPS II? Other than the wiki page that is in the Trac system, there is nothing specific, though this question has been asked before.

*Update to action: Jordan Gerth will follow-up to see if there is further documentation. John Olsen may know of some documentation out there.*

Local models in GRIB2 formatting dictate you need to have a center of origin. We would like a different center of origin that does not lead to encroachment of the different data header values that redefine the grid origin and record type. There has been some traffic on this over awips2dev but will try to follow up with John Olsen.

- **Obtain and distribute the latest documentation on localization.**
  - **OPR: Ed Mandel (complete by Oct 29)**

There have been updates to the localization document in December. Ed said they were now prototyping some automation.

*Update to action: Ed Mandel will take the action to get the latest documentation and see the automated scripts, as well as the automation documentation as well.*

- **GOES-R Proving Ground Partners: In preparation for the next TIM:**
  - **Submit any specific questions to be considered at next TIM.**
  - **Gather specific topics based on the results of what came out of this TIM.**
  - **Submit items which each participant on the call is working on or concerned about.**
  - **Submit items which require a design-oriented discussion. Think long term.**
  - **See action items above.**

The date for the in-person TIM is unknown. We are aiming for the summer. The GRPG all-hands in-person meeting will occur in the middle of May. There may be a telephonic TIM at that time.

Frank Griffith will be at AMS. Raytheon will have an AWIPS II live system with data at their large booth. NCEP is going to be part of the booth too, so you will see nmap.

*An ongoing action item is to submit questions for consideration at these TIMs.*